

HP ProCurve Switches

HP ProCurve Switch 5300xl Series

Date 2/11/2003

Version 1.1



HP ProCurve Switch 5372xl



HP ProCurve Switch 5348xl

Table of Contents

1. BACKGROUND.....	5
1.1 HP ProCurve Networking	5
1.2 Hewlett-Packard 5300 Switch Products	5
1.2.1 HP ProCurve Switches Covered in this Guide	5
1.3 HP ProCurve Adaptive EDGE Architecture™	6
1.4 HP Switch Positioning.....	7
1.4.1 Positioning for the HP ProCurve Switch 5308xl.....	8
1.4.2 Positioning for the HP ProCurve Switch 5304xl.....	8
2. EVALUATION FEATURES AND BENEFITS	9
2.1 Feature Set Summary.....	9
2.1.1 Architecture	9
2.1.2 High Availability	9
2.1.3 Prioritization / QoS.....	9
2.1.4 Security.....	9
2.1.5 Bandwidth Management.....	10
2.1.6 Network Management	10
2.1.7 Availability.....	10
2.1.8 Service and Support.....	10
2.1.9 New Features in Software Release E.07.0x.....	11
2.2 Architecture	11
2.2.1 Hardware Architecture Summary	11
2.2.2 N-Chip.....	11
2.2.2.1 Classification and Lookup.....	12
2.2.2.2 N-Chip Programmability	12
2.2.2.3 Fabric Interface.....	12
2.2.2.4 The N-Chip CPU	12
2.2.3 F-Chip.....	13
2.2.4 The Master CPU	13
2.3 High Availability.....	13
2.3.1 IP Routing.....	13
2.3.2 Rapid Spanning Tree Protocol, 802.1w.....	14
2.3.2.1 IEEE 802.1D Spanning Tree Protocol.....	14
2.3.3 Switch Meshing (LAN Aggregation)	15
2.3.4 XRRP – Router Redundancy Protocol	16
2.4 Prioritization / QoS	17
2.4.1 Priority Queues.....	17
2.4.2 QoS Classifiers	17
2.4.3 IEEE 802.1p Priority Support.....	18
2.4.4 Diffserv / TOS Support	18
2.4.5 End-to-End QoS	19
2.5 Security.....	19
2.5.1 Filtering	19
2.5.1.1 ACLs – Access Control Lists	19

2.5.1.2	Static Filters.....	20
2.5.2	802.1x – Port-based access control / RADIUS Authentication	20
2.5.2.1	RADIUS Server Accounting	21
2.5.2.2	Standalone RADIUS Authentication.....	21
2.5.2.3	RADIUS Functionality - RFCs.....	21
2.5.3	TACACS+ Authentication	22
2.5.4	Port Security - MAC Lockdown.....	22
2.5.5	Secure Shell – SSHv2.....	22
2.5.6	SSL – Secure Sockets Layer.....	23
2.5.7	Management VLAN	23
2.5.8	SNMPv3	23
2.5.9	Manager Authorized List.....	23
2.6	Bandwidth Management.....	24
2.6.1	Port Trunking – (Port Aggregation)	24
2.6.1.1	802.3ad – LACP	24
2.6.1.2	Cisco Fast EtherChannel®	24
2.6.1.3	Trunking in a Layer 3 Environment	24
2.6.2	VLANs	25
2.6.2.1	IEEE 802.1Q VLAN Support.....	25
2.6.2.2	GVRP	25
2.6.3	IGMP	26
2.7	Network Management.....	26
2.7.1	MIB Support.....	26
2.7.2	RMON Support	27
2.7.3	Network Monitoring Port	27
2.7.4	Console Support	27
2.8	Availability.....	28
2.8.1	Hot Swap	28
2.8.2	Redundant Power Supply	28
2.8.3	Dual Flash.....	28
2.8.4	Alert Log	28
2.8.4.1	SysLog Capabilities.....	29
2.9	Service and Support	29
2.9.1	Lifetime Software Updates (Best in the Industry)	29
2.9.2	Lifetime Warranty (Best in the Industry)	29
2.9.3	Telephone Support.....	29
2.9.4	Optional Support Services	29
3.	PERFORMANCE.....	31
3.1	IP Routing (L3) RFC 2285 Fully Meshed Throughput Test.....	31
3.1.1	Copper Gigabit ports	31
3.1.2	100BT Ports.....	31
3.1.3	Throughput Test Comments	31
3.2	IP Routing (L3) RFC 2245 Latency Test.....	31
3.2.1	Copper Gig Ports	31
3.2.2	100BT Ports.....	32
3.2.3	Latency Test Comments	32
3.3	5300 vs. the Cisco Catalyst 4006 Tolly Report.....	32
4.	ADDITIONAL INFORMATION.....	33

4.1 ProCurve Networking Web Site.....33

5. PRICING35

1. Background

This HP ProCurve Switch 5300xl Series Reviewer's Guide will help network engineers at computer trade publications, resellers and end-user sites evaluate the merits of the HP ProCurve switches.

1.1 HP ProCurve Networking

HP ProCurve Networking is an extensible line of products working together to provide the control network administrators need to deliver the network uptime and performance that their organizations require. This guide describes one part of HP networking: the HP ProCurve Switch 5300xl Series products.

1.2 Hewlett-Packard 5300 Switch Products

1.2.1 HP ProCurve Switches Covered in this Guide

This guide covers the following Hewlett-Packard switch products:

HP ProCurve Switch 5308xl (J4819A)

The HP ProCurve Switch 5308xl is a 5U wire speed Layer 2/3/4 eight slot chassis switch targeted primarily at high performance access tier applications. The HP ProCurve Switch 5308xl provides up to 32 Gigabit Ethernet ports or 192 10/100-TX with HP Auto-MDI-X ports. Voice, video and data ready, the Switch 5300xl series offers extensive prioritization that bring full convergence down to the desktop. The chassis comes standard with a routing engine and power supply.

HP ProCurve Switch 5304xl (J4850A)

The HP ProCurve Switch 5304xl is a 3U wire speed Layer 2/3/4 four slot chassis switch targeted primarily at high performance access tier applications. The HP ProCurve Switch 5304xl provides up to 16 Gigabit Ethernet ports or 96 10/100-TX with HP Auto-MDI-X ports. Voice, video and data ready, the Switch 5300xl series offers extensive prioritization features that brings full convergence down to the desktop. The chassis comes standard with a routing engine and power supply.

HP ProCurve Switch 5372xl (J4848A)

The HP ProCurve Switch 5372xl bundles 72 10/100-TX ports in the 8 slot 5308xl chassis, leaving 5 open slots.

HP ProCurve Switch 5348xl (J4849A)

The HP ProCurve Switch 5348xl bundles 48 10/100-TX ports in the 4 slot 5304xl chassis, leaving 2 open slots.

Modules available for these switches:

- HP ProCurve Switch xl 10/100Base-TX module (J4820A)
24 ports of 10/100Base-T
- HP ProCurve Switch xl 100/1000-T module (J4821A)
4 ports of 100/1000Base-T (no 10Mb support)
- HP ProCurve Switch xl 100FX MR-RJ module (J4852A)
12 ports of 100FX – MT-RJ connectors
- HP ProCurve Switch xl mini-GBIC module (J4878A)
4 ports of mini-GBIC connectivity
 - HP ProCurve Gigabit-SX-LC Mini-GBIC (J4858A)
 - HP ProCurve Gigabit-LX-LC Mini-GBIC (J4859A)
 - HP ProCurve Gigabit-LH-LC Mini-GBIC (J4860A)

- HP ProCurve Switch redundant power supply (J4839A)

Pricing for each of these switches and modules is given in the Pricing section.

1.3 HP ProCurve Adaptive EDGE Architecture™

Networks are now being asked to carry many different types of data with differing delivery requirements. Traffic volumes are rising quickly. Security needs are more stringent. And the typical user is demanding a higher level of mobility than ever before. To meet these critical needs HP ProCurve Networking has developed the HP ProCurve Adaptive EDGE Architecture. The two major tenets of the EDGE Architecture are:

- Intelligent control to the edge, and
- Command from the center

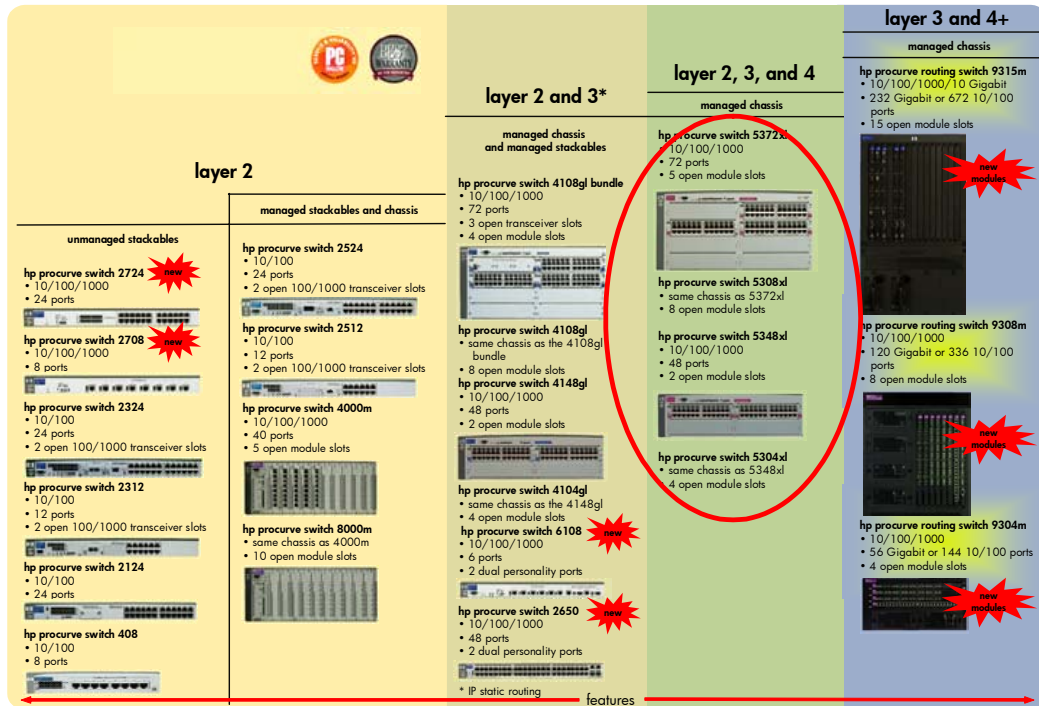
It is the network edge where users and applications connect, where network traffic enters and exits the network, and where the network must determine how that traffic should be handled. The edge is where security policies must be enforced, where the user connects after being authenticated at a central command resource. Without control to the edge, decisions about security and traffic must be deferred to the network core, impacting core performance and scalability while at the same time requiring more bandwidth in all parts of the network driving up cost and complexity. In addition, this opens the network to security attacks between where access is physically attained and where authorization is granted. The intelligent control to the edge must be done in the switches closest to the users. Since these switches constitute the highest number of network ports in a network, they must also be cost effective.

The Adaptive EDGE Architecture is not just a future vision. Many elements of the architecture are already available in HP ProCurve's current products, including the HP ProCurve Switch 5300 Series. With its HP developed ASICs, the 5300 series delivers a broad range of Layer 2, 3 and 4 features for control to the edge. The 5300 series is cost effective at the edge and can be coupled with the 9300 series core switches or used to create a distributed core in a network comprised entirely of 5300 series switches in a meshed – at layer 2 or layer 3 – highly available configuration or grid.

Over time more features will be added to the HP ProCurve Switch 5300 Series to round out the intelligent control to the edge, filling in the command from the center, providing a dynamic network environment needed by users in a rapidly evolving information environment.

For more information on the HP ProCurve Adaptive EDGE Architecture, see the HP ProCurve website at <http://www.hp.com/go/hpprocurve>.

HP ProCurve product portfolio



1.4 HP Switch Positioning

A widely used method for segmenting the areas in which switches are installed calls for three different classifications: access, distribution and core. Access switches provide aggregation of end nodes for connection to a distribution or core switch and are usually found in wiring closets. Distribution switches aggregate the links from access switches and possibly server farms. Distribution switches anchor the network in a building, or for small networks, across a campus. Core switches provide the focal point of the local network, aggregating the distribution switches, providing connectivity for central site data centers, and providing connectivity in many cases to the external network.

Access switch requirements vary depending on the use model of the customer. Some customers just want basic aggregation with high speed uplinks. Other customers require a more sophisticated approach with security, QoS, Layer 3 routing services, VLAN services, flexible filtering, and some level of fault tolerance. The HP ProCurve Switch 5300xl Series meets the needs of the sophisticated access tier implementation.

Other switches in the HP ProCurve switch product line cover the needs of the lower end access tier, as well as distribution and core areas of the network.

1.4.1 Positioning for the HP ProCurve Switch 5308xl

The HP ProCurve Switch 5308xl is a 5U wire speed Layer 2/3/4 eight slot chassis switch targeted primarily at high performance access tier applications where higher port density is needed. The HP ProCurve Switch 5308xl provides up to 32 Gigabit Ethernet ports or 192 10/100-TX with HP Auto-MDI-X ports. Voice, video and data ready, the Switch 5300xl series offers extensive prioritization that bring full convergence down to the desktop. The chassis comes standard with a routing engine and power supply.

The HP ProCurve Switch 5308xl is the same as the HP ProCurve Switch 5304xl, except that it holds up to 8 modules in a 5U rack space, giving it a higher port density and greater throughput than the HP ProCurve Switch 5304xl. For some customers the HP ProCurve Switch 5308xl can also be used as a distribution or core switch.

1.4.2 Positioning for the HP ProCurve Switch 5304xl

The HP ProCurve Switch 5304xl is a 3U wire speed Layer 2/3/4 four slot chassis switch targeted primarily at high performance access tier applications. The HP ProCurve Switch 5304xl provides up to 16 Gigabit Ethernet ports or 96 10/100-TX with HP Auto-MDI-X ports. Voice, video and data ready, the Switch 5300xl series offers extensive prioritization features that bring full convergence down to the desktop. The chassis comes standard with a routing engine and power supply.

The HP ProCurve Switch 5304xl can be used in wiring closets or server farms where high performance at Layer 2 or Layer 3, particularly at Gigabit rates, is desired. The 3U chassis size lends itself to smaller racks or where the larger number of ports available through the HP ProCurve Switch 5308xl is not needed. For some customers the HP ProCurve Switch 5304xl can also be used as a distribution or core switch.

2. Evaluation Features and Benefits

2.1 Feature Set Summary

The HP ProCurve Switch 5300xl Series are store-and-forward Layer 2/3/4 routing switches. Features, discussed in more detail later in this section, include:

2.1.1 Architecture

- A high speed Layer 3 architecture consisting of a full routing switch ASIC on each module, all interconnected via the backplane crossbar switch fabric ASIC. Both ASICs are HP designed. Up to 16,536 (16K) L2 MAC addresses are supported.

2.1.2 High Availability

- IP Functionality supported:
 - Routing services: RIP (v1, v1 compatible v2, and v2), OSPF, static routes
 - 10,000 network address routes, 65,536 (64K) L3 host address routes
 - IPv4 routing, IPv6 switching
 - 16 multi-netted interfaces per VLAN
 - DHCP relay – allows DHCP requests to be forwarded to links associated with the DHCP server
- IEEE 802.1w Rapid Spanning Tree Protocol support – provides very fast Spanning Tree convergence (approaching 1 second under optimal conditions) on lost links or when the root switch is unreachable. Compatible with switches running 802.1D Spanning Tree.
- XRRP Router Redundancy Protocol: Two 5300s can back each other up for Layer 3 interfaces. Failure detection and switch-over can be as fast as 3 seconds.
- HP Layer 2 Switch Meshing: Allows fully meshed connections between switches at Layer 2 with all links being used to send traffic.

2.1.3 Prioritization / QoS

- Four priority queues
- Traffic prioritization based on:
 - UDP/TCP Application Type (port number)
 - Device Priority (destination or source IP address)
 - IP Type of Service (ToS/Diffserv) field (IP packets only)
 - Protocol Priority (IP, IPX, ARP, DEC LAT, AppleTalk, SNA, and NetBEUI)
 - VLAN Priority
 - Incoming source-port on the switch
 - Incoming 802.1p Priority (present in tagged VLAN environments)

2.1.4 Security

- Filtering capabilities include Access Control Lists (ACLs), source port, multicast MAC address and protocol filtering capabilities.
- 802.1x – port based access control
- RADIUS / TACACS+ authentication

- Port security – MAC lockdown
- SSHv2 – secure connection for telnet services
- SSL – Secure Sockets Layer for secure interaction between a browser and the 5300's management GUI interface
- Management VLAN – Limit CLI/GUI/telnet access to the switch to a particular VLAN
- Manager authorized list – limiting access to the Web services, telnet and tftp

2.1.5 Bandwidth Management

- 802.3ad LACP (Port Trunks) – (Link Aggregation Control Protocol) Switch-to-switch and switch-to-server aggregated links allow scalable bandwidth communication. Can be used in many cases to trunk to non-HP devices. Also supported is Cisco's Fast EtherChannel® with PAgP.
- 256 VLANs, 802.1Q compliant, Layer 2 port-based, for segmentation of local LANs
- Support of GVRP (part of 802.1Q) for automatic configuration of VLANs throughout a Layer 2 environment
- IGMPv3, IGMP snooping, data-driven IGMP

2.1.6 Network Management

- Web-based management for anytime, anywhere configuration access
- HP Tootools for Hubs & Switches (standalone)
- HP Tootools For Hubs & Switches integration into HP OpenView-NT, HP OpenView-UX, CA Unicenter, and Tivoli via no-cost bridge software between these platforms and HP Tootools.
- HP OpenView/NT native application via the HP ProCurve network management for OV-NT product
- HP OpenView/UX native application via the HP ProCurve network management for OV-UX product
- Alert Log capability which finds common network problems and informs the net manager of the situation
- SNMPv3/MIB II/RMON support on all ports for encrypted monitoring and control
- Ability to configure a network monitoring port for use with external probes or analyzers
- HP's Extended RMON support allowing the monitoring of traffic flows in the network

2.1.7 Availability

- Hot swap capability, load-sharing power supplies, dual flash memory

2.1.8 Service and Support

- Lifetime warranty (for as long as you own the product) with next business day advanced replacement (available in most countries)
- Free lifetime software updates
- Free telephone support during business hours.
- Optional fee-based services, such as upgrading your warranty to on-site response, or 24X7 phone support.

2.1.9 New Features in Software Release E.07.21 (released January 22, 2003)

- ACLs
- SSHv2
- IGMPv3
- Debug/Logging
- XRRP
- SSL
- Meshing improvements
- 100FX module software support
- SNMPv3
- OSPF Route Authentication

2.2 Architecture

2.2.1 Hardware Architecture Summary

The HP ProCurve Switch 5304xl has 4 identical slots, while the HP ProCurve Switch 5308xl has eight. Any of the Switch 5300xl modules can be put in any of the slots.

The switch architecture is based on 2 different HP designed ASICs: the Network or N-Chip, and the Fabric or F-Chip. Each module has an N-Chip that provides on-module routing and switching functions. It also provides the high speed connection to the backplane. The F-Chip, located on the backplane, provides the wire speed crossbar fabric interconnecting all the modules. This combination of highly integrated N-Chips connected through the F-Chip gives the HP ProCurve Switch 5300xl Series the ability to deliver wire-speed Layer 3 for the price of Layer 2 switching, and in a chassis form factor.

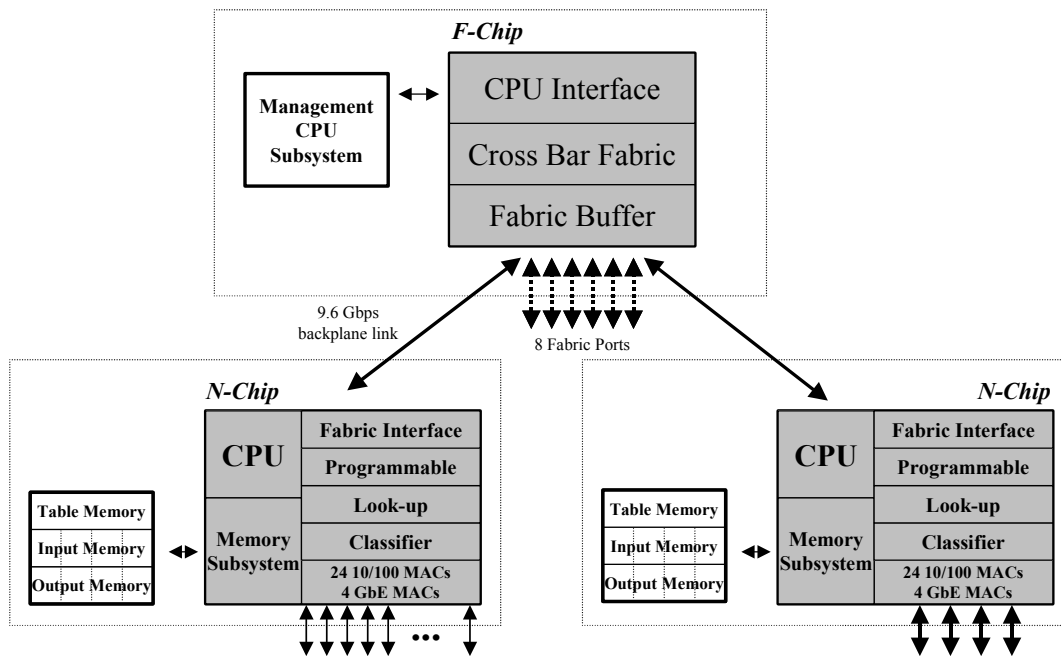


Figure 1. Detailed Architecture

The HP ProCurve Switch 5300xl Series have two slots in the back for the load-sharing power supplies. One power supply ships standard with each switch and can power a fully loaded chassis. A second power supply can be installed for redundancy and longer overall expected power supply life.

The HP ProCurve Switch 5300xl Series can hold up to 16,536 (16K) MAC addresses in the switch address table.

2.2.2 N-Chip

Each module contains a full ASIC-based Layer 3 routing switch engine. This switch engine, called the network or N-Chip, provides all the packet processing: Layer 2 and Layer 3 lookups, filtering and forwarding decisions, VLAN, trunking and priority queuing determinations. The N-Chip also contains its own CPU.

2.2.2.1 Classification and Lookup

When a packet first comes in, the classifier section determines the packet characteristics, its addresses, VLAN affiliation, any priority specification, etc. The packet is stored in input memory, lookups into the table memory are done to determine routing information and a N-Chip specific packet header is created for this packet with this information. This header is then forwarded to the programmable section of the N-Chip.

2.2.2.2 N-Chip Programmability

As mentioned in the previous section, one of the functions of the N-Chip is to analyze each packet's header as it comes into the switch. The packet's addresses can be read, with the switch making forwarding decisions based on this analysis. For example, if a packet's 802.1Q tag needs to be changed to re-map the packet priority, the N-Chip needs to look at each packet to see if any particular one needs to be changed. This packet-by-packet processing has to occur very quickly to maintain overall wire-speed performance. ASICs (application specific integrated circuits) provide this high performance, but typically cannot be changed in their functionality once the ASIC design is frozen.

To broaden the flexibility of the N-Chip, a programmable function is included in some areas of its packet processing. This programmability provides network processor-like capability, giving the HP designers the opportunity to make some future changes or additions in the packet processing features of the ASIC by downloading new software into it. Thus new features needing high performance ASIC processing can be accommodated, extending the useful life of the switch without the need to upgrade or replace the hardware.

This programmable functionality was originally designed and implemented in the popular HP ProCurve Switch 4000M switch family and was used to give the HP ProCurve Switch 4000M new ASIC-related features well after initial release of the product. Customers with existing units could benefit from the new features via a free software download. The customer's investment in the Switch 4000M was preserved by providing new functionality not otherwise possible without the ASIC programmability.

Being based on the Switch 4000M's implementation, the HP ProCurve 5300xl programmable capability is a second generation design.

2.2.2.3 Fabric Interface

After the packet header leaves the programmable section, the header is forwarded to the Fabric Interface. The Fabric Interface makes final adjustments to the header based on priority information, multicast grouping, etc. and then uses this header to modify the actual packet header as necessary.

The Fabric Interface then negotiates with the destination N-Chip for outbound packet buffer space. If congestion on the outbound port is present, WRED (weighted random early detection) can also be applied at this point as a congestion avoidance mechanism.

Finally the N-Chip Fabric Interface forwards the entire packet through the F-Chip to an awaiting output buffer on the N-Chip that controls the outbound port for the packet. Packet transfer from the N-Chip to the F-Chip is provided via the 9.6Gbps full duplex backplane connection, also managed by the Fabric Interface.

2.2.2.4 The N-Chip CPU

The N-Chip contains its own CPU, a 66 MHz ARM-7, for Layer 2 learns, packet sampling for the XRMON function, handling local MIB counters and running other module related operations.

Overall, the local CPU offloads the master CPU by providing a distributed approach to general housekeeping tasks associated with every packet. MIB variables, which need to be updated with each packet, can be done locally. The Layer 2 forwarding table is kept fresh via this CPU. Other per-port protocols, such as Spanning Tree, LACP and CDP, are also run on this CPU.

The local CPU, being a full-function microprocessor, allows functionality updates through future software releases.

2.2.3 F-Chip

The fabric, or F-Chip, which is located on the backplane of the switch, provides the crossbar fabric for interconnecting the modules together. The use of a crossbar allows wire speed connections simultaneously from any module to any other module. As mentioned in the N-Chip section, the connection between the F-Chip and each N-Chip (module) in the chassis is through a 9.6Gbps full duplex link.

One unique function of the F-Chip is to automatically replicate multicast packets and send them to the destination modules. This method is more efficient than having the source N-Chip do the replication. Since only a single copy of the multicast packet needs to be sent to the F-Chip, this method saves bandwidth on the high speed connection between the source N-Chip and the F-Chip.

2.2.4 The Master CPU

Along with the F-Chip, the backplane of the switch also contains the master CPU, 32MB RAM and 12MB of flash ROM memory. The master CPU, a 200 MHz Power PC 8240, runs the routing protocols and maintains the master routing tables, maintains the master MIBs, responds to SNMP requests, and manages the user interfaces. The Master CPU is also responsible for switch bootup coordination. Two copies of the switch operating system can be stored in the flash ROM. This allows the user to recover quickly if the main code copy is corrupted or a code update produces results other than what is desired.

Input to the CPU is prioritized into 4 queues. Queuing this way prevents the user from being locked out of the switch user interface due to unintentional high levels of traffic, such as broadcast storms. More significantly, this also prevents a user lockout due to intentionally high levels of traffic, such as denial of service attacks.

2.3 High Availability

2.3.1 IP Routing

IP routing on the HP ProCurve Switch 5300xl Series is done in the ASIC at wire speed by the user defining VLANs and then specifying routing between them. Some of the IP services available are:

- Routing Services
 - RIP (version 1, version 1 compatible version 2, and version 2)
 - Split Horizon and Poison Reverse supported
 - Redistribution - importing of static and connected routes into the RIP route table. Restrict command available to prevent route advertising in and/or out of any port. Useful for security (routed sections of the network can be made invisible to the rest of the network.)
 - Up to 128 routed interfaces. Since RIP sends out its full routing table on each routed interface once every 30 seconds, care should be taken to limit the number of interfaces used for larger routed environments. If the number of routes needed is greater than 1,000 with a higher number of routed interfaces, for instance, greater than 32, then OSPF should be considered as the routing protocol of choice, as it is more efficient in handling larger routed environments.
 - OSPF (RFC 1583 (default) and 2328 compliant)
 - Redistribution – importing of static and connected routes into the OSPF route table. Restrict command available to prevent route advertising in and/or out of any port. Useful for security (routed sections of the network can be made invisible to the rest of the network.)
 - OSPF traps (RFC 1850)
 - OSPF Route Authentication – uses plain text passwords or MD5 encryption. The MD5 keys are time sensitive – the 5300 must have the correct time/date set, particularly after a reboot. It is recommended that when using OSPF route authentication TimeP or SNTP (time setting protocols) be used to assure a proper time/date setting, particularly on reboot caused by a power failure.
 - 128 routing interfaces (interfaces that participate in the routing protocols) per chassis

- Static IP routes: 512 maximum per chassis
- IRDP (ICMP Router Discovery Protocol)
- Proxy ARP
- Up to 10,000 network address routes – enough for a large local environment
- Up to 65,536 Layer 3 host address routes
- Bootp Relay Service
- Encapsulation type: Ethernet II
- 8 Subnets per VLAN: one primary subnet and up to 7 secondary subnets. Maximum 512 secondary subnets per chassis. (Max 256 primary subnets (VLANs) + 512 secondary subnets = 768 max total subnets per chassis)
- DHCP relay – allows DHCP requests to be forwarded across routed interfaces to links associated with the DHCP server
- IPv4 routing, IPv6 switching: full routing of IPv4-based packets, IPv6-based packets are switched at Layer 2

2.3.2 Rapid Spanning Tree Protocol, 802.1w

Spanning Tree Protocol (STP), part of the IEEE 802.1D standard, is a Layer 2 protocol that allows switches to be interconnected with redundant multiple links for high availability that form network loops. In a non-spanning tree environment these loops would immediately bring the network down. Using link cost algorithms, Spanning Tree determines which redundant links should be logically shut down thus preventing any active network path loops.

There are two concerns with the original Spanning Tree standard, 802.1D. The first is that all redundant links except one are not used for actual network traffic. This wastes potential bandwidth. This problem is usually addressed by routing at the switch instead of just switching. Many network managers don't want to do this, however, due to the higher level of management needed in a routed environment over a Layer 2 environment. While the HP ProCurve Switch 5300xl Series can perform Layer 3 routing, there is an easier solution in Switch Meshing, which is described in the next section.

The second concern is on link failure or loss of the STP root switch, Spanning Tree can take up to 45 seconds to re-establish network connections. In many networks a potential outage of 45 seconds is unacceptable. While many switch vendors in the past have implemented a proprietary protocol to deal with this, the IEEE has now established the 802.1w standard, Rapid Spanning Tree Protocol (RSTP) to update the STP so that it responds more rapidly to link failure or loss of the root switch. Actual recovery time is dependent (as STP is) on network complexity but can approach 1 second under optimal conditions. RSTP is better than the proprietary protocols because it is standards based, leading to interoperability between different switch vendors, and it provides backward compatibility with the original STP. Sections of the network that are running under STP will recover with times associated with STP, while those running under RSTP recover in RSTP timeframes.

2.3.2.1 IEEE 802.1D Spanning Tree Protocol

The HP ProCurve Switch 5300xl Series also support single instance spanning tree, per the 802.1D specification. Running STP and RSTP in the same box is mutually exclusive; only one form can be run at any given time. RSTP is the recommended configuration and can be run in the same spanning tree domain with other switches that are running STP. STP is available, however, for users that for some reason don't want to run RSTP.

2.3.3 Switch Meshing (LAN Aggregation)

The HP ProCurve Switch 5300xl Series family supports HP's Switch Meshing, a way to interconnect these switches in a meshed topology at Layer 2. Meshed switch-to-switch links can all be used simultaneously to their full advantage, with traffic being load-balanced through redundant links based on dynamically determined latency on the different possible paths between switches. Highly available, fault tolerant networks can be easily built with very low network administration required.




Figure 2. Switch Meshing

Note, however, that routing and Switch Meshing cannot be used in the same switch at the same time.

In traditional switched environments, meshed topologies are not allowed without the use of the Rapid or original Spanning Tree Protocol (802.1w or IEEE 802.1D). Spanning Tree detects loops in the topology and logically blocks as many links as necessary to avoid traffic loops. If one of the active links fails, Spanning Tree enables another link to re-establish the path, if possible. Unfortunately, Spanning Tree requires links to be available that are not being used for data, letting available bandwidth go unused.

Although RSTP and STP are supported by the HP ProCurve Switch 5300xl Series, Switch Meshing is superior since all available links are used between switches. With Switch Meshing, the switch selects the best traffic path for each new destination end-node it learns based on dynamically determined latency on each of the possible paths to the node. Recalculation of path latency in each switch is done every 30 seconds and is based on link speeds, input and output buffer queue lengths, and knowledge of any dropped packets on particular ports.

Redundancy is also provided by Switch Meshing. If one of the links fails between switches, traffic is redirected through another path, if available. The switchover time of typically less than 1 second is very fast. Very robust high availability solutions can be implemented with a switch mesh.

Switch Meshing allows multiple HP ProCurve Switch 5300xl Series to form a virtual backplane between the switches, allowing reliable high port density environments to be made inexpensively.

Up to 12 switches can participate in a Switch Meshing domain, with up to 5 switch hops between the most distant switches in the mesh. Multiple Switch Meshing domains can exist in a single LAN environment, but not within the same switch.

Routing switches and routers use a similar technique through routing protocols such as RIP or OSPF. In many situations Switch Meshing is an improvement over these routing protocols because:

- The path decision in HP's Switch Meshing is made using dynamically determined latency through the switches. Routing protocols do not take latency into account, only path costs based on link speeds (OSPF) or simply the lowest number of router hops (RIP).
- Switch Meshing works for all Layer 3 protocols, as well as non-routable protocols such as DEC LAT or NET BIOS, because path specification is performed using Layer 2 MAC addresses. Routing specifies paths based on supported Layer 3 protocols (usually IP, sometimes IPX and rarely AppleTalk), otherwise the router must simply bridge the packet and use Spanning Tree.
- Configuration of Switch Meshing is trivial. Specifying which ports are part of the Switch Meshing domain is all that is needed. The switch takes care of the rest. This is in sharp contrast to configuration of routing protocols which can be challenging.
- Convergence time (time to recover from a lost link) is fast - typically less than one second. This is much faster than RIP and faster or on par with OSPF using triggered updates.

- Unlike a router, no packet modification is required as it travels through the switch.

Other HP ProCurve switches that support Switch Meshing, such as the HP ProCurve Switch 4000M family, will work together in a mesh with the HP ProCurve Switch 5300xl Series. There are a few corner case caveats in this type of mixed environment that are covered in the HP ProCurve 5300xl documentation. The HP ProCurve 5300xl documentation can be found at: <http://www.hp.com/go/hpprocurve> under the Technical Support section.

A white paper with more details on Switch Meshing can be found in the information library on HP's networking web site at <http://www.hp.com/go/hpprocurve>.

2.3.4 XRRP – Router Redundancy Protocol

One form of high availability in a Layer 3 environment is having two routing switches back each other up. In the event of a connection failure with one of the routing switches, the other routing switch transparently takes over the routing function. XRRP, the XL Router Redundancy Protocol, provides the mechanism in the HP ProCurve Switch 5300xl Series routing switches for this backup functionality.

Similar in concept to VRRP (Virtual Router Redundancy Protocol), XRRP presents a virtual router to the end node connections whose IP and MAC address can transition from the master HP ProCurve 5300 to the backup HP ProCurve 5300 on master 5300 interface failure. Since the end node connections are tied to the virtual router using the virtual router IP and MAC addresses, they are unaware as to whether the actual physical routing services are being provided by the master 5300, or, after a switch-over, to the backup 5300, making any switch-over transparent to the end nodes. An XRRP interface failure is defined as the inability of the master physical interface in the 5300 pair to be heard by the backup interface. This could be caused by a cable failure, module failure, whole 5300 failure, or operator error (such as a disconnected cable).

Some XRRP specifications:

- Number of physical routers in a backup group (XRRP calls this a 'protection domain'): 2
- Number of protection domains allowed per VLAN: 16
- Time to failure detection and switchover: default – 15 seconds. Minimum time by making a configuration change – 3 seconds. If a VLAN is lost on one of the 5300 pairs, but the 5300 doesn't go down, fail-over occurs in under 1 second as the 5300 with the failed VLAN reports the loss directly to the other 5300 via a different VLAN.
- Backed-up interfaces should be configured identically between the routers. XRRP checks and warns if interface configurations do not match.
- Master interfaces can be split between the two 5300 switches, allowing a split of traffic between the two 5300s under normal network conditions.
- If a failure is detected on any master interface, all the XRRP master interfaces on that 5300 are switched over to the back-up router. This allows easier troubleshooting, or module or box replacement. When all interfaces on the failed 5300 are restored, the master relationship is re-established as it was before the fail-over. There is a time interval (XRRP fail back – default 10 seconds, configurable to 999 seconds) before master re-establishment can take place to prevent master interface flapping due to interfaces that may be going up and down.
- Those interfaces not defined as part of the XRRP set on a 5300 will continue to run on that 5300 (unless, of course, the whole 5300 is down) even as the XRRP interfaces switch over to the backup 5300. This has value for low priority interfaces where the cost of redundant resources across two 5300s for these interfaces is not cost justified.

- If a 5300 Management VLAN is enabled it cannot be defined as an XRRP interface – SNMP management requests to a particular 5300 need to go to that physical switch regardless of fail-over status.
- XRRP does not interoperate with VRRP, but can coexist in a VRRP environment without interference.

For more details see the HP ProCurve Switch 5300xl Series documentation located at: <http://www.hp.com/go/hpprocurve> under the Technical Support section.

2.4 Prioritization / QoS

Quality of Service (QoS) mechanisms in the HP ProCurve Switch 5300xl Series provide the network manager control over packet flows based on a number of factors. In addition, since the switches can override the priority values in the incoming packets, the network manager can maintain QoS control over inappropriate priority designations coming from users or applications at the end nodes. Conversely, many applications can be given priority treatment through the switch without the end node clients having to be aware of QoS, particularly valuable since client operating systems and the applications themselves are generally not QoS aware at this time.

The primary means of control is through priority queues in the switch. Pieces of information in the packet that can be used to determine priority queue placement are called classifiers. The mechanism to actually store the priority based on the classifiers is through the 802.1Q tags or through the IP TOS/Diffserv section in the IP header. The HP ProCurve Switch 5300xl Series do not modify either of the packet fields when routing the packet (unless a QoS override is specified in the port config) so the QoS status of a packet can be maintained as the packet travels elsewhere in the network. Each of these is discussed in the following sections.

The final section, End-to-End QoS, briefly discusses the value of QoS in networks.

2.4.1 Priority Queues

Each port on an HP ProCurve Switch 5300xl Series module has four priority queues. A packet placed into a particular queue will be processed according to the priority of that queue.

The priority queues are managed through a fair-weighted queuing algorithm that prevents any priority queue from getting starved (the packets in it not being serviced by the switch) due to high traffic levels in higher priority queues.

2.4.2 QoS Classifiers

Through user configuration, priority of packets can be specified based on the following classifiers. This list is in order of precedence; if there are multiple classifiers that apply to a specific packet, the one that is highest on this list takes effect.

- Layer 4 TCP/UDP port numbers: allows prioritization based on the application associated with the packet. This allows, for instance, VoIP packets using fixed port numbers to be prioritized higher than other traffic. It can also be used to downgrade packet flows, such as HTTP traffic. Can also be used to remap the diffserv code points (DSCP).
- Device Priority (destination or source IP address) : up to 256 addresses can be specified per chassis, destination address takes precedence over source address. Can also be used to remap the DSCP.
- IP Type of Service (ToS) field (IP packets only): support for both the older TOS IP Precedence definition, or the newer Differentiated Services (Diffserv) definition. If using the TOS IP Precedence, the bits are mapped to packet priority queues using the same table as shown in the next section “IEEE 802.1p Priority Support”. The 802.1p bits are also set for the outbound packet if the packet goes out of the switch through a port

that has 802.1Q tagging turned on.

For diffserv, each of the diffserv code points (DSCP) can have a priority set for it. It is also possible to set a new DSCP and 802.1p priority based on the incoming DSCP, or set the 802.1p priority alone based on the incoming DSCP. The ability to re-write the DSCP allows the network manager to:

- Identify packets coming from a different area of the network, such as a remote site, by changing the DSCP as it comes through the HP ProCurve Switch 5300xl Series and treating this remote packet differently than packets originating in the local environment
- Redefine an incoming DSCP to conform to the DSCP definitions defined in the local environment.
- Protocol Priority (IP, IPX, ARP, DEC LAT, AppleTalk, SNA, and NetBEUI)
- VLAN ID: Allows a VLAN to be assigned a specific priority. Can also remap the DSCP.
- Incoming source-port on the switch. Can also remap the DSCP.
- Incoming 802.1p Priority (present in tagged VLAN environments)

If the DSCP is remapped, the 802.1p priority associated with the new DSCP is used to determine the priority queue on the outbound port. In addition, these 802.1p bits will be included in the outgoing packet if VLAN tagging is specified for the outbound port.

2.4.3 IEEE 802.1p Priority Support

IEEE 802.1Q packet tagging supports both designation of VLAN membership (see the VLAN section below) and packet priority (up to 8 levels and often referred to as 802.1p). Since 802.1p has 8 levels of priority possible, but the switch has only 4 physical priority levels, the following mappings are used:

802.1p priority	Switch priority queue
1, 2	1 (low)
0, 3	2 (normal)
4, 5	3
6, 7	4 (highest)

Packets without any 802.1p tagging are assigned by the switch internally to 802.1p priority 0. This is mapped to the normal queue in the switch so that untagged packets are not penalized in priority. These priority queue mappings are set as designated in the 802.1Q standard.¹

2.4.4 Diffserv / TOS Support

As mentioned in the Classifiers section above, the HP ProCurve Switch 5300xl Series provide very flexible control of the diffserv (DSCP) bits. Mapping of each of the 64 possible DSCPs can specify an 802.1p priority, as well as a new DSCP for the outbound packet.

Diffserv code points are in their early acceptance for use in networks. Their importance will grow as more networks and applications take advantage of them.

Since TOS IP Precedence and DSCP are mutually exclusive (they use the same set of bits in the IP header), the switch will allow only TOS IP Precedence definitions or only DSCP definitions to be active at any one time.

¹ Some other switch vendors use non-standard priority mappings in their switches.

2.4.5 End-to-End QoS

QoS capabilities in the switch allow it to deal with two different concerns that arise in Ethernet networks: congestion control and latency. In the past, controlling traffic congestion was viewed as the primary reason for QoS. But with Ethernet prices dropping substantially year after year it has been easier and lower cost to deal with congestion, at least in the local LAN, by increasing the bandwidth available to traffic through higher speed connections.

While QoS for congestion control in the local LAN has had marginal value, the ability of QoS to deal with applications that are sensitive to varying latencies through a network is of value. Delay sensitive applications depend on isochronous, or time-dependent, data. Applications of this type include VoIP, streaming voice or video, data storage backups, or network control in the form of SNMP packets, Spanning Tree BDPUs, etc.

When trying to make overall packet latency as low as possible or minimize latency jitter, end-to-end control becomes important. The 802.1p priority specification that is contained in each tagged packet, as well as the DSCP, can provide this end-to-end continuity. As the packet travels from source to destination, it is given the proper priority in each switch it travels through based on its 802.1p value. The HP ProCurve Switch 5300xl Series maintain the 802.1p tags across routed interfaces, allowing end-to-end QoS in routed environments.

The DSCP can also be used for QoS categorization of the packet. The HP ProCurve Switch 5300xl Series can assign priorities based on the DSCP. Packets that are not 802.1p tagged can nonetheless have a priority assigned to them through the DSCP alone.

The ability of the HP ProCurve Switch 5300xl Series to control not only the 802.1p priority, but also read and/or re-write the DSCP bits to set QoS policy provides the network manager with an even finer degree of control. Priority can be tailored to specific areas of the network, and the DSCPs can be used eventually for parts of QoS policy other than priority. There is also room in the DSCP definition for new QoS services that have not yet been defined.

2.5 Security

2.5.1 Filtering

2.5.1.1 ACLs – Access Control Lists

When routing is turned on across Layer 3 interfaces, all routable packets are allowed across these interfaces. Selectively filtering the packets that can flow across these interfaces is useful for security or bandwidth control purposes. Filtering at Layer 3 is done through ACLs, Access Control Lists.

A single complete filter statement, the ACL, is comprised of one or more ACEs, Access Control Entries. An ACE statement can permit or deny a packet based on it's:

- Source and/or destination IP address or IP subnet
- Source and/or destination TCP/UDP port number with less than, greater than, equal, not equal or number range. Being able to specify less than, greater than, etc. can save a lot of ACEs trying to bound a group of port numbers and is not found in some competitors' ACL implementations.
- IP protocol (IP, TCP, UDP)

Each static VLAN on the 5300 can have one inbound and one outbound ACL defined. The 5300 can have up to 99 Standard ACLs, which are defined as ACLs that are based only on source IP addresses. The 5300 can also have up to 99 Extended ACLs, which are defined as ACLs based on any of the other parameters listed above. Up to a total of 1024 ACEs can be used to specify the 5300 ACLs.

The order of the ACEs within the ACL is important. When processing an ACL, the 5300 starts with the first ACE in the ACL and will continue to work through the list of ACEs, in order, until the packet matches the condition set forth in a particular ACE. At that point any further ACEs in the ACL are

ignored. If a packet does not match any of the conditions in the ACL, it is denied. This is in keeping with the typical use of ACLs as a security mechanism. If the automatic denial property is not wanted, the ACL should end with an ACE statement permitting ANY. To assist in writing and editing ACLs, the ACL file can be edited externally and downloaded into the 5300.

A typical use for standard ACLs is to allow a single end node on one subnet access to a server on another subnet, while denying all other ends nodes on the first subnet similar access. An example of this situation would be an human resource representative getting access to a personnel database on another subnet, while keeping all other end nodes from accessing this same database. Similarly, a Standard ACL could be used to deny access of an entire subnet to anywhere in the corporate network other than out to the Internet.

Extended ACLs can be used as filters for application traffic that uses fixed TCP/UDP port numbers. For example, an Extended ACL can be set up to only allow traffic from a particular subnet access to the email servers on another subnet. Or an extended ACL could deny any traffic destined for custom applications (those applications using port numbers above 1024).

The ACL functionality of the HP ProCurve Switch 5300xl Series supports ACL logging. When logging is specified in a particular ACE, an entry is made in the log when that ACE results in an explicitly denied packet. Logging of permitted packets is not supported. The 5300 ACL logging is primarily useful for troubleshooting.

ACLs, being a Layer 3 service in the 5300, are only executed for packets that are routed, crossing a VLAN/router boundary. They have no effect on packets that are being switched in a Layer 2 environment.

ACLs for the HP ProCurve Switch 5300xl Series are flexible and can be used to create sophisticated filters. Before implementing ACLs, ACL details should be consulted in the HP ProCurve Switch 5300xl Series documentation located at: <http://www.hp.com/go/hpprocurve> under the Technical Support section.

2.5.1.2 Static Filters

Static filtering can be used to provide security and/or bandwidth control within the network. When a static filter is defined it can be applied to any or all ports on the switch. The following three types of static filters can be defined:

- Source port: Packets coming from a particular port can be dropped. Source port filters can be used to isolate ports from each other and allow communication only to uplinks, for example. Ports that can use a particular source port filter must be in the same VLAN as the source port. Up to 78 source port filters can be defined on the chassis
- Multicast MAC address: If an IGMP group is active in the address range of a static multicast filter, IGMP takes precedence. Once the IGMP group becomes inactive, the static multicast filter takes affect. Up to 16 multicast address filters can be defined
- Protocol type: up to 7 protocol filters. Protocols that apply to the protocol filter are:
 - AppleTalk
 - DEC LAT
 - NetBEUI
 - ARP
 - IP
 - IPX
 - SNA

These filters are done in hardware; there is no performance penalty when using them.

2.5.2 802.1x – Port-based access control / RADIUS Authentication

The IEEE 802.1x standard governs a methodology for client system network log-in. Through 802.1x a user is given access to the network only after the HP ProCurve Switch 5300xl Series (the network access server) authenticates the user through a RADIUS server. As part of this authentication, the user

can be given specific network access rights, such as assignment to a specific VLAN and some high level session accounting information can be maintained. (See the next section.)

With a centralized RADIUS server doing the actual authentication, a user can log-in anywhere in the network that supports 802.1x and get access to his resources. This is true whether the log-in occurs on a shared client, or the user is using a mobile client and accessing the network at different access points.

One point to note about 802.1x: access control is that it is control to the port of the switch. Once access is given to the switch port, anyone connected through this port will have access to the services associated with the user that authenticated. If someone inadvertently, or clandestinely places a switch or hub between the network access server and the authenticated client, any port on the introduced switch or hub has access to the configured network services of the authenticated client. One way to close this shortfall is to use the Port Security MAC Address Lockdown feature on the HP ProCurve Switch 5300xl Series, which is described in a following section.

More details on 802.1x can be found in the white paper on the HP ProCurve website at <http://www.hp.com/go/hpprocurve> (select the information library).

2.5.2.1 RADIUS Server Accounting

Most RADIUS servers can provide not only authentication for the user, but can also keep track of some parameters associated with the authenticated user or the switch itself. These parameters are actually kept on the HP ProCurve Switch 5300xl Series and updated on the RADIUS server at either RADIUS session begin/end or just at session end.

Three areas of parameters are tracked:

- Network Accounting – Keeps track of items for an authenticated user on a switch port such as Account ID, Username, Input and Output Packets, Account Termination Reason, etc.
- Exec Accounting – Keeps track of the same items used in Network Accounting, but for logon sessions under telnet, SSH and console.
- System Accounting – Keeps track of the same items used in Network Accounting, with actual recording of the items done on a system event, such as system reboot, system reset and accounting enable or disable

The primary purpose for RADIUS accounting is to have a security audit trail for user network usage or when switch events occur that affect the integrity of the network.

RADIUS server accounting can also be used as a rudimentary form of tracking user network usage, but only covers very high level parameters such as total connect time, or total packets through the user's switch port.

2.5.2.2 Standalone RADIUS Authentication

RADIUS authentication can be used without using 802.1x. In this case RADIUS is used to provide user authentication when telnet, SSH or console port access authentication is required. Up to three RADIUS servers can be specified to provide backup capability in case the primary RADIUS server becomes unavailable.

2.5.2.3 RADIUS Functionality - RFCs

RFCs that were used or consulted in the development of the RADIUS functionality are:

- RFC-2865 - Remote Authentication Dial In User Service (RADIUS)
- RFC-2869 - RADIUS Extensions
- RFC-2138 - Extensible Authentication Protocol Support in RADIUS
- draft-congdon-radius-8021x-09.txt - IEEE 802.1X RADIUS Usage Guidelines

- RFC-2868 - RADIUS Attributes for Tunnel Protocol Support
- RFC-2618 - RADIUS Authentication Client MIB
- RFC-2866 - RADIUS Accounting
- RFC-2620 - RADIUS Accounting Client MIB

2.5.3 TACACS+ Authentication

The HP ProCurve Switch 5300xl Series supports TACACS+ as an authentication means for switch telnet or console port access. The switches support two levels of access: if the user/password combination listed on the TACACS+ server is given a privilege level of 15 the user has Manager access (read/write) to the switch. A privilege level of 14 or lower will restrict the user to Operator status (read only).

Backup TACACS+ servers can be configured providing multiple TACACS+ server access in case the primary TACACS+ server is unavailable for any reason.

2.5.4 Port Security - MAC Lockdown

The 802.1x standard provides logical security to the network based on a user. There are many times, however, when physical access limitations are desired. The Port Security - MAC Lockdown feature limits physical access to a particular port on the switch by one of two methods: a particular list of MAC addresses (up to 8 addresses per port can be configured), or to the first MAC address the switch sees on that port. While this solution doesn't help with a switch port that legitimately sees a large number of MAC addresses, such as in a conference room, it does provide security to a port used by a shared PC or dedicated PC by locking out other PCs that try to access the switch port, even when the port is network enabled through 802.1x.

The Port Security feature can be set to send an SNMP trap to a management station when such a violation occurs. It can also be set to completely disable the switch port (requiring the network manager to re-enable the port before use), a feature for use in high security environments, or an environment subject to potential hacking, such as a college dorm room.

2.5.5 Secure Shell – SSHv2

Secure Shell is an application very similar to telnet except that it encrypts the dialog so that in-band CLI sessions can be kept private over the network. Encryption is done through the use of public/private key pairs, one pair for host authentication and one pair for each SSH session that is initiated.

The host key pair is used to authenticate the SSH client and switch to each other. The host key pair is stored in flash, so is not lost on reboot, power-cycle or by clearing the config file. Although not necessary or recommended, a new host key pair can be generated through the CLI.

The session key pair is used to authenticate the SSH session. A new key pair is used for each SSH session. Keys are kept in RAM and are lost on power-cycle or reboot. When the HP ProCurve Switch 5300xl Series is rebooted, new session key pairs are generated. With a key pair taking about 12 seconds to generate, 10 keys are generated on boot up and placed in a cache to prevent delays when starting up SSH sessions rapidly in succession. Filling this key cache takes about 2 minutes and is CPU intensive. To keep this process from affecting other switch functions, it is designated low priority for the CPU. Because the CPU is doing many things at boot up, key pair generation doesn't start until about one minute after boot up. This means that an SSH session, waiting for the first session key pair generation, cannot be established until a little over a minute after boot up.

The HP ProCurve Switch 5300xl Series support both SSHv1 and SSHv2 clients. SSHv2 provides an additional level of security in that the public key negotiation is accomplished via a Diffie-Hellman exchange that is not done under SSHv1.

2.5.6 SSL – Secure Sockets Layer

SSL can be used to encrypt the exchange between a web browser and the 5300 switch when using the HP ProCurve Switch 5300xl Series web GUI.

A facility is provided on the GUI interface to generate a self-signed RSA certificate for use during a SSL browser session.

2.5.7 Management VLAN

The HP ProCurve Switch 5300xl Series can be configured to designate one of the VLANs to be the management VLAN. When this is configured the internal IP address of the switch becomes a member solely of the management VLAN. Since access to the switch IP address is necessary for telnet/SSH, GUI, and SNMP access, other members of this VLAN are the only ones that can manage the switch.

The management VLAN is useful when higher switch security is desired. It prevents general switch function access by anyone other than those on the management VLAN. The management VLAN cannot be designated an XRRP backup VLAN.

2.5.8 SNMPv3

Many functions of the HP ProCurve Switch 5300xl Series can be monitored and the switch configuration can even be changed through the switch's MIBs. The standard method of querying the switch's MIBs for network management is through SNMP, the simple network management protocol.

Before version 3 of SNMP, SNMP has used clear text across the network. On some networks this has been viewed as a possible serious security concern. A way around this has been to use a network management specific VLAN (see the section above on Management VLAN), but this can be restrictive and is not a viable solution in many environments, particularly remote environments.

SNMPv3 provides security for the SNMP communications across the web, including an encryption mechanism to encrypt packet information. The three levels of security available in SNMPv3 are:

- Authentication between the SNMP initiator and the 5300 switch based on username. Not very secure.
- Authentication between the SNMP initiator and the 5300 switch based on MD5 or SHA algorithms. Better security for the passwords as they are encrypted. Actual SNMP communication after login is still clear text and not secure.
- Authentication between the SNMP initiator and the 5300 switch based on MD5 or SHA algorithms and encryption via 56 bit key DES. Passwords are protected and further SNMP communication is encrypted across the network. Querying and control via SNMP cannot be viewed outside the encrypted session.

With SNMPv3 those sites that are concerned with the possibility of packet snooping can turn on encryption allowing secure communication between the network management application and the switch.

2.5.9 Manager Authorized List

The HP ProCurve Switch 5300xl Series Manager Authorized List can be configured with up to ten IP addresses that have management access to the switch. The list, along with Management VLANs and console passwords, provides a way to tightly limit who has access to the switch console.

If no addresses are in this list (the default) any source IP address can send a packet to the switch's management agent. If you do have addresses in this list and you are using a management VLAN, addresses on the list must be a member of the management VLAN to obtain switch login.

2.6 Bandwidth Management

2.6.1 Port Trunking – (Port Aggregation)

Link Aggregation is the industry term for the ability to combine multiple coterminous links (links that begin at the same point and end at the same point) as one logical link.

Link aggregation allows two HP ProCurve switches to be interconnected by 2-4 of the same type of links, with all links acting as one higher-speed link. Since the number of links in a trunk is configurable, the bandwidth is scaleable to the needs of a particular network. For example, 4 links at 100Mbps can be trunked to provide the equivalent of a 400 Mbps (800 Mbps full-duplex) link between two switches or 4 individual Gigabit links can be trunked for the equivalent of a 4 Gigabit (8 Gigabit full-duplex) link. Fiber-optic links can be trunked to interconnect switches across large campuses. Port trunking also provides redundancy on links between the two switches or switch and server. If one of the links fails, the traffic is moved to another link in the trunk in under one second.

The HP ProCurve Switch 5300xl Series support 36 port trunks of up to 4 physical links each. There are 3 ways to configure which ports on the switch participate in trunks: LACP (802.3ad), Cisco Fast EtherChannel®, and manually. Although an easy process, manually configured trunks do require the user to configure them directly into the switch. Any changes in the links used will require a manual change in this configuration. The advantage of manual configuration is that it allows the HP ProCurve Switch 5300xl Series to work with trunks from a number of other vendors that do not adhere to the LACP standard or support Fast EtherChannel®.

2.6.1.1 802.3ad – LACP

Automatic configuration of port trunks happens when using LACP, Link Aggregation Control Protocol. LACP, in the active form, operates by sending out packets looking for LACP running on the other end of each connection. The switch, however, has a default configuration of passive LACP²; each link is listening for an active LACP connection on the other end. For LACP to dynamically configure the ports in a trunk, one or both ends of the trunks need to have LACP configured in its active mode. Once the user configures active LACP, links can be moved to different ports, or new links can be added, with LACP detecting this and reconfiguring to reestablish the LACP trunk.

LACP, like the other forms of trunking, supports 4 links per trunk. LACP does allow the configuration of standby links. Standby links carry no data unless one of the active links in the trunk fails. Standby links are used in situations where the loss of even one of the active links would cause an unacceptable traffic load on the remaining active links.

2.6.1.2 Cisco Fast EtherChannel®

Another form of port trunk configuration protocol is Cisco's proprietary Fast EtherChannel® and Gigabit EtherChannel®. Under EtherChannel®, the control protocol is called PAGP, Port Aggregation Protocol. It essentially performs the same function as LACP but is not standards based. Since many Cisco products do not currently support LACP, PAGP can be used to automatically configure the trunk links between the HP ProCurve Switch 5300xl Series and Cisco switches.

2.6.1.3 Trunking in a Layer 3 Environment

Traditional trunking uses MAC (Layer 2) addresses to determine which link in the trunk a particular traffic flow travels over to avoid the problem of out-of-sequence packets. In a Layer 3 environment between two routing switches this would cause all packets to flow over only one link because the source and destination MAC addresses for all packets would be the same – the MAC address of the two connected routing switches.

² The passive form of LACP is the default because active LACP sends out packets periodically on each port looking for LACP on the other end. While this traffic level is very low, most users don't want any traffic that doesn't directly relate to their environment.

To avoid this situation the HP ProCurve Switch 5300 Series uses the source and destination IP addresses to determine which link a particular packet flow uses. This will provide a good overall distribution of traffic across the different links in the trunk.

2.6.2 VLANs

A Virtual LAN is a logical collection of ports or nodes that belong to a single broadcast/multicast domain. VLANs were originally devised as a solution to limit the size of any one broadcast domain to allow scaling of switched environments. With the advent of routing switch solutions, however, use of VLANs in end user environments is now largely done for network policy or security reasons.

For the HP ProCurve Switch 5300xl Series, VLANs are also used to provide entities to which to attach the router functionality. All routing in the HP ProCurve Switch 5300xl Series is defined to be between VLANs.

HP ProCurve Switch 5300xl Series support 256 VLANs (8 default). VLAN membership can be designated through either a particular port (untagged), or through a 802.1Q tag.

VLANs can overlap on a single port. For example, it may be advantageous to have a server connected through a single port to be a member of two different VLANs³ such that two different groups of PCs can access the same server, but the two groups of PCs cannot talk directly with each other.

2.6.2.1 IEEE 802.1Q VLAN Support

The HP ProCurve Switch 5300xl Series support the IEEE 802.1Q VLAN tagging standard. The HP ProCurve Switch 5300xl Series can have multiple VLAN traffic streams share a single physical link. 802.1Q also allows interoperability at this level between different vendors in a standards-based way. End-to-end VLAN designation is also greatly simplified through the 802.1Q tag, particularly if GVRP, discussed in the next section, is used.

Ports with only a single VLAN designation can be designated as untagged ports. Packets leaving these ports will not be 802.1Q tagged. VLAN continuity from switch-to-switch must be manually maintained at each switch if untagged ports are used.

2.6.2.2 GVRP

GVRP—GARP VLAN Registration Protocol is a standard under 802.1Q that provides a facility to dynamically configure a VLAN on switches throughout a Layer 2 domain when that VLAN has been statically configured on at least one switch in the domain. The intention with GVRP is to automatically interconnect a VLAN that is manually configured on two switches that are not contiguous in a Layer 2 domain. This greatly reduces the administrative overhead of having to define VLANs in all the intermediate switches between two VLAN islands that need to be interconnected. GVRP will also delete a dynamic VLAN on any switch port that hasn't heard externally from the VLAN in the last 10 seconds.

GVRP is particularly advantageous in environments using 802.1x, network login. In 802.1x, as a user is authenticated to the switch from the RADIUS server, a VLAN membership can also be indicated. This allows the network manager to assign a particular user to a particular VLAN to establish the network services available for that user. For example, when the user logs in they can be placed in their own VLAN along with the servers and storage that contains the services that user is allowed to have. One of the advantages of 802.1x is the ability for the user to login anywhere in the network. If they are assigned to a VLAN that also has services members, that VLAN will have to be defined all along the path between the user and those services. GVRP will automatically do this. GVRP will also delete that VLAN along the path once it is no longer needed.

The HP ProCurve Switch 5300xl Series has a configuration default of 8 VLANs maximum that can be defined. If GVRP is enabled, the 'maximum VLANs to support' value should be configured in most

³ If multiple VLANs to a server are done using 802.1Q, the server must also support 802.1Q tagging.

cases to be the maximum number of VLANs expected in the entire network environment supported by GVRP. This will provide space for dynamic VLAN definitions that come about through GVRP.

2.6.3 IGMP

Internet Group Multicast Protocol (IGMP) is a multicast control protocol that builds delivery paths through the switch network. The switch eavesdrops on IGMP traffic so that it knows which ports are part of which multicast groups. If there is no multicast router available, the HP ProCurve Switch 5300xl Series can act as an IGMP querier to learn which end nodes have subscribed to which multicast streams. The switch can then direct a specific multicast stream to only those switch segments that have nodes that have joined the multicast group associated with that stream. Configuration of this feature is a single check box to turn it on. The HP ProCurve Switch 5300xl Series supports IGMP version 3 by recognizing and processing IGMPv3 joins. Version 3 is backwards compatible with versions 1 and 2.

The HP ProCurve Switch 5300xl Series support a maximum of 389 IGMP groups.

2.7 Network Management

Network management is an important part of a network solution. There are three levels of network management available for the HP ProCurve Switch 5300xl Series:

- Web-based management - Configuration of the individual switch can be done anytime, anywhere through the web server available in each switch, accessible via a standard web browser.
- HP Tootools for Hubs & Switches - Management of a network of HP hubs, switches and routing switches can be done through the included HP Tootools for Hubs & Switches application, part of the HP Tootools suite of management products. HP Tootools for Hubs & Switches provides a network map and device configuration in a web-based format.
- HP OpenView platforms - If management of a generic (HP and other vendors' devices) or large (>1500 managed nodes) IP network is needed:

NT platform – HP ProCurve Network Management for OV-NT (J4869B) can be separately purchased to provide direct management of the HP ProCurve Switch 5300xl Series and other HP ProCurve products through HP OpenView/NT. For more details see the product information at <http://www.hp.com/go/hpprocurve>.

HP-UX platform - Management of the HP ProCurve Switch 5300xl Series under HP OpenView/UX is available via the separately purchased J3250P HP Hub & Switch Management for OV-UX product.

TopTools itself will also run under HP Openview, CA Unicenter, and Tivoli using the proper no-cost bridges to these environments. See the TopTools web site at <http://www.hp.com/toptools> for more details.

Up to SNMP version 3 is supported. See the security details of SNMPv3 in the Security section above.

2.7.1 MIB Support

The HP ProCurve Switch 5300xl Series supports the following standard MIBs:

- MIB-II (RFC 1213)
- Ether-like MIB (RFC 1398)
- Bridge MIB (RFC 1493)
- Ethernet-like MIB (RFC 1643)
- RIP v2 MIB (RFC 1724)
- OSPF MIB (RFC 1850)
- RMON probe configuration MIB – RMON v2 (RFC 2021)

- IP Forwarding Table MIB (RFC 2096)
- SMON MIB (RFC 2613)
- RADIUS Client MIB (RFC 2618)
- RADIUS Client Accounting MIB (RFC 2620)
- Ethernet MIB (RFC 2665)
- 802.3 MAU MIB (RFC 2668)
- 802.1p and 802.1Q Bridge MIB (RFC 2674)
- Entity MIB (RFC 2737)
- RMON MIB (RFC 2819)
 - Groups: (1) Ethernet statistics, (2) Ethernet history, (3) Alarm, (9) Event,
- Evolution of Interface MIB (RFC 2863)

In addition, a number of enterprise-specific MIBs are also supported for such things as VLANs, XRRP, and multiple bridge groups.

2.7.2 RMON Support

For those customers that use RMON applications, the HP ProCurve Switch 5300xl Series support RMON groups 1 (Ethernet statistics), 2 (Ethernet history), 3 (Alarm), and 9 (Event). These four groups are available for all ports.

The Ethernet statistics group provides counters for packet counts, broadcast/multicast packets, packet length counts, missed packets and erred packets. Event and Alarm groups allow threshold setting and alarm generation based on the counters in the Ethernet group. History accumulates two records for each port, one a sampling of Ethernet statistics taken every 30 seconds and the other a sampling of statistics taken every 30 minutes.

2.7.3 Network Monitoring Port

If more RMON groups (such as packet capture) are desired, a RMON probe can be attached to one of the HP ProCurve 5300xl ports and Port Monitoring can be configured. Port Monitoring allows the end user to copy all traffic, inbound and outbound, from any number of ports, even those on different subnets within the switch, to a single destination port. This allows the probe to see all traffic on the selected ports to provide the probe with the proper global perspective.

The Network Monitoring Port can also copy all traffic for one particular VLAN to the destination port, rather than having to specify on a port-by-port basis.

This feature is very helpful when using a LAN analyzer when doing specific monitoring or troubleshooting of network segments.

2.7.4 Console Support

Out-of-band management of the HP ProCurve Switch 5300xl Series can be done through the RS-232 console port via a directly connected terminal emulator. The console interface provides three ways to configure the switch:

- Setup – provides a quick, simple one screen menu to set up the switch for items such as IP address, time protocol method, community name, spanning tree, etc. without having to use the command line interface. Particularly useful for getting an IP address into the switch so that additional configuration can be done through the browser GUI.
- Menu – provides easy menu style configuration and monitoring of the major areas of the switch. Many switch configurations can be entirely done through the menu system without any real knowledge of the command line interface.
- CLI (command line interface) – provides configuration and monitoring access to every function on the switch.

The console interface is also available in-band through the network via the telnet service.

The RS-232 port of the switch speed senses the RS-232 port on the terminal/PC interface up to 115,200 baud. Modems are also supported through the RS-232 port using a straight through cable.

2.8 Availability

Availability is the measure of the ability for the switch to remain running over a period of time with minimal impact to the network environment.

2.8.1 Hot Swap

The HP ProCurve Switch 5300xl Series allows hot-swapping of the port modules and hot-swapping of the mini-GBIC while the switch is still running without affecting the other port modules. This allows a defective port module to be replaced without affecting the rest of the network.

2.8.2 Redundant Power Supply

An optional redundant power supply (identical to the primary supply) can be installed. The power supplies load-share, allowing both supplies to run cooler, extending their Mean Time Between Failure (MTBF) values. If either of the supplies in the chassis does fail, the other continues to run preventing switch interruption. Power supplies can be hot-swapped (unplug it first!) during switch operation when two are present. Each supply has its own power cord and it is recommended that they be connected to different power mains to limit exposure to failure in any one power main.

2.8.3 Dual Flash

The HP ProCurve Switch 5300xl Series have dual flash memory. This provides for two copies of the switch operating system and is particularly useful when doing an OS upgrade. If problems are found when moving to the new OS, the switch can be immediately rebooted using the older OS.

2.8.4 Alert Log

The HP ProCurve Switch 5300xl Series, like most other switches in the HP ProCurve line, look for the following common port-based network problems:

- Too many undersized/giant packets
- Excessive CRC/alignment errors
- High collision or drop rate
- Loss of link
- Excessive jabbering
- Excessive late collisions
- Excessive broadcasts

When any of these conditions are detected on a switch port, the HP ProCurve Switch 5300xl Series informs the network manager through:

- The browser-based GUI interface (Alert Log on the switch status page). The GUI interface will also provide some suggested remedies for these problems when the user double-clicks on the alert in the list.
- A SNMP trap sent to the net management application(s) address(es) configured on the switch

This alert feature can save a lot of troubleshooting time, particularly in small network environments that may not be continuously running a network management application. Many of these detected problems tend to be intermittent and thus difficult to troubleshoot. Having a readily available port specific alert log, with time stamps and possible problem remedies, can speed up troubleshooting resolution, particularly for a very part-time network manager.

In larger environments, proactive messages sent to the net management station speeds detection and can lower troubleshooting time, as alarm thresholds do not have to be set specifically in the net management application to capture these problems.

2.8.4.1 SysLog Capabilities

As described in the previous section, local logging is a feature well suited for a small network. When a network becomes much larger than a few switches, having to go to each individual switch makes diagnosing a problem more difficult. Syslog functionality, growing out of the UNIX environment, is a method of sending log entries to a centralized server. This enables system and network administrators to analyze and troubleshoot their entire network from a central location.

The HP ProCurve Switch 5300 Series can send all log entries to a specific server (via an IP address), or send varying severity levels of log entries to different servers. Severity levels supported are:

- emergency
- alert
- critical
- error
- warning
- notice
- info
- debug

2.9 Service and Support

Hewlett-Packard has long been known for its high quality products and excellent service and support of them. HP ProCurve switches are no exception.

2.9.1 Lifetime Software Updates (Best in the Industry)

As with other HP ProCurve products, the HP ProCurve Switch 5300xl Series come with lifetime software updates. For as long as software updates are available for these switches they can be downloaded from the HP web site for free. This is an industry-leading product feature.

2.9.2 Lifetime Warranty (Best in the Industry)

Warranties, and the ease of obtaining warranty service for the end user, is a product benefit that is easily overlooked in a technical evaluation, but ranks high as a concern of end users as they get ready to actually make a purchase decision. The HP ProCurve Switch 5300xl Series have a limited lifetime warranty, for as long as you own the product. If any part of the switch fails due to a defect in material or workmanship, including the power supply or fans, it will be replaced. In most parts of the world, the replacement unit is sent with next business day delivery in advance of the failing unit being returned to HP. Advance replacement gets the unit to the end user as fast as possible and minimizes down time by allowing the impaired unit to continue to be used if possible until the replacement unit arrives. This also allows for easy scheduling for when the actual unit swap occurs on the network. An optional upgrade to on-site replacement is also available in most parts of the world. Refer to the warranty statement that ships with the product for more details on warranty coverage.

The HP ProCurve Switch 5300xl Series warranty is industry-leading.

2.9.3 Telephone Support

HP provides free pre-sales and post-sales telephone support during normal business hours to end users and HP resellers through the HP Customer Care Centers located world-wide.

2.9.4 Optional Support Services

In addition to free support services such as the warranty and telephone support, HP offers an extensive range of fee-based support services to meet more specialized needs. The following optional services are available for the HP ProCurve Switch 5300xl Series:

- Onsite next business day
- Onsite in 4 hours same business day
- Onsite in 4 hours 24x7
- 6 hour call-to-repair
- 24x7 telephone support

Hewlett-Packard can also provide more broad-based services such as site surveys, installation services, and actual management of the network, depending on customer needs.

More information can be found at <http://www.hp.com/go/hpprocurve> or by contacting a local HP sales office.

3. Performance

These numbers have been generated by Hewlett-Packard, using testers from Ixia Communications. Ixia testers are used by a number of network testing houses and the press to determine performance numbers for networking equipment. In these tests, 32 ports were used for Gigabit testing, 192 ports for 100 Mb testing. All ports were full duplex. Numbers presented here are condensed from Ixia reports in order to save space.

Testing done on the HP ProCurve Switch 5308xl. Maximum rate of throughput (100%) would be the same for the 5304xl but at one-half the number of packets since the 5304xl has one-half the possible number of ports of the 5308xl.

3.1 IP Routing (L3) RFC 2285 Fully Meshed Throughput Test

3.1.1 Copper Gigabit ports

Port pairs active, full duplex: 32 = 32 Gbps data out of the tester

Test length: 5 minutes

Packet size (bytes)	64	128	256	512	1024	1280	1518
%MaxRate	100	100	100	100	100	100	100
TotalTxFrames	14285711648	8108112000	4347829856	2255634400	1149426432	923077824	780229824
TotalRxFrames	14285711648	8108112000	4347829856	2255634400	1149426432	923077824	780229824
TotalLoss(%)	0.000	0.000	0.000	0.000	0.000	0.000	0.000

3.1.2 100BT Ports

Port pairs active, full duplex: 192 = 19.2 Gbps data out of the tester

Test length: 3 hours

Packet size (bytes)	64	128	256	512	1024	1280	1518
%MaxRate	100	100	100	100	100	100	100
TotalTxFrames	308572144115	175134177024	93913320982	48721238434	24827090688	19937539584	16852104192
TotalRxFrames	308572144115	175134177024	93913320982	48721238434	24827090688	19937539584	16852104192
TotalLoss(%)	0.000	0.000	0.000	0.000	0.000	0.000	0.000

3.1.3 Throughput Test Comments

A fully meshed performance test sends packets from each port to every other port during the test. This test exercises both the modules and the backplane. These tests show the HP ProCurve Switch 5300xl Series to be wire-speed on all ports simultaneously. The 5300 is the only chassis in its price range that is wire-speed on all ports simultaneously at Layer 2 or Layer 3.

3.2 IP Routing (L3) RFC 2245 Latency Test

3.2.1 Copper Gig Ports

Port pairs active, full duplex: 32

All latencies in microseconds

Frame Size	64	128	256	512	1024	1280	1518
AvgLatency(μs)	5.802	6.746	8.085	11.731	18.067	21.104	23.940

3.2.2 100BT Ports

Port pairs active, full duplex: 192

All latencies in microseconds

Frame Size	64	128	256	512	1024	1280	1518
AvgLatency(μs)	24.36	36.26	42.38	81.44	136.46	166.42	200.82

3.2.3 Latency Test Comments

Latency is measured as the time it takes for a byte inside a packet to enter and then leave the switch. This measurement includes both the processing time of the switch as it makes its forwarding decision and the time for the packet itself to enter and leave the switch.

The latency figures for the HP ProCurve Switch 5300xl Series are low. Latencies this low will not be a factor in general network operation, even with streaming video or VoIP applications.

Almost all switches currently on the market, the 5300 included, are store and forward, so the entire packet is received into the switch before the switch starts to transmit it out the outgoing port. The above latency figures include this packet receive time. For example, at 100Mbps it takes 5.76 μsec for a 64 byte packet, and 122.08 μsec for a 1518 byte packet itself to move into and out of the switch. At 1Gbps a 64 byte packet takes 576 nseconds, while a 1518 byte packet takes 12.208 μsec. Adding the packet receive time to the latency is proper because this extra time is seen externally to the switch by the network and figures in to the transit time of the packet as it moves through the network.

3.3 5300 vs. the Cisco Catalyst 4006 Tolly Report

HP ProCurve asked The Tolly Group, a well-known and respected test and consultancy firm, to compare the performance of the HP ProCurve Switch 5300xl Series against the Cisco Catalyst 4006. Highlights of this testing are:

- Delivers 100% of wire speed Layer 2 throughput with Gigabit Ethernet uplinks compared to the Catalyst 4006 which delivered only 25% in the same scenario
- Exhibits at least 20% lower latency for all packet sizes
- Conforms more closely to defined mapping and management of multiple priority queues than the Catalyst 4006
- Offers four times better performance at one-fourth of the price of the Catalyst 4006 with Gigabit-over-copper ports

The entire Tolly Group report is available on the HP ProCurve web site at:

<http://www.hp.com/go/hpprocurve>

4. Additional Information

4.1 ProCurve Networking Web Site

Additional information, including the latest data sheets, design services, white papers, product documentation and support information can be obtained through the HP ProCurve Networking web site. HP ProCurve Networking can be reached at:

<http://www.hp.com/go/hpprocurve>

The information contained in this document is subject to change without notice.

5. Pricing

All managed HP switches ship with HP TopTools for Hubs & Switches.

Prod No.	Description	US List Price February 1, 2003
J4819A	HP ProCurve Switch 5308xl Chassis with 1 power supply, routing engine, and 8 open module slots	\$2,999
J4848A	HP ProCurve Switch 5372xl Switch 5308xl pre-configured with 72 10/100 ports. Includes 1 power supply, routing engine, and 5 open module slots	\$7,129
J4850A	HP ProCurve Switch 5304xl Chassis with 1 power supply, routing engine, and 4 open module slots	\$1,999
J4849A	HP ProCurve Switch 5348xl Switch 5304xl pre-configured with 48 10/100 ports. Includes 1 power supply, routing engine, and 2 open module slots	\$4,759
Modules		
J4820A	HP ProCurve Switch 10/100Base-TX module 24 autosensing 10/100 RJ45 ports	\$2,379
J4821A	HP ProCurve Switch 100/1000Base-T module 4 autosensing 100/1000 RJ45 ports	\$1,099
J4852A	HP ProCurve Switch 100FX MT-RJ module 12 ports of 100FX – MT-RJ connectors	\$4,729
J4878A	HP ProCurve Switch Mini-GBIC module 4 port module for Gigabit Ethernet mini-GBIC (SFP: SX, LX)	\$1,299
Mini-GBICs		
J4858A	HP ProCurve Gigabit-SX-LC mini-GBIC	\$479
J4859A	HP ProCurve Gigabit-LX-LC mini-GBIC	\$1,059
J4860A	HP ProCurve Gigabit-LH-LC mini-GBIC	\$5,699
RPS:		
J4839A	HP ProCurve Switch Redundant Power Supply	\$1,099
Network Management		
J4869B	HP ProCurve Network Management for OpenView-NT	\$3,699
J3250P	HP Hub & Switch Management for OV-UX	\$2,119

Pricing is subject to change without notice.